BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images, please do not report the images to the Image Problem Mailbox.

Terms used **coherency** or **coherence** and **directory** and **owner** or **ownership** and **tag**

Found **2,272** of **138,517**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Adjustable block size coherent caches

Czarek Dubnicki, Thomas J. LeBlanc

April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual international symposium on Computer architecture**, Volume 20 Issue 2

Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms

Several studies have shown that the performance of coherent caches depends on the relationship between the granularity of sharing and locality exhibited by the program and the cache block size. Large cache blocks exploit processor and spatial locality, but may cause unnecessary cache invalidations due to false sharing. Small cache blocks can reduce the number of cache invalidations, but increase the nuber of bus or network transactions required to load data into the cache. In this paper we ...

2 Using destination-set prediction to improve the latency/bandwidth tradeoff in shared-memory multiprocessors

Milo M. K. Martin, Pacia J. Harper, Daniel J. Sorin, Mark D. Hill, David A. Wood

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2

Full text available: pdf(220.76 KB) Additional Information: full citation, abstract, references

Destination-set prediction can improve the latency/bandwidth tradeoff in shared-memory multiprocessors. The destination set is the collection of processors that receive a particular coherence request. Snooping protocols send requests to the maximal destination set (i.e., all processors), reducing latency for cache-to-cache misses at the expense of increased traffic. Directory protocols send requests to the minimal destination set, reducing bandwidth at the expense of an indirection through the d ...

3 Piranha: a scalable architecture based on single-chip multiprocessing

Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture**, Volume 28 Issue 2

Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

4 Memory coherence in shared virtual memory systems

Kai Li, Paul Hudak

November 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 4

Full text available: pdf(2.71 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The memory coherence problem in designing and implementing a shared virtual memory on loosely coupled multiprocessors is studied in depth. Two classes of algorithms, centralized and distributed, for solving the problem are presented. A prototype shared virtual memory on an Apollo ring based on these algorithms has been implemented. Both theoretical and practical results show that the memory coherence problem can indeed be solved efficiently on a loosely coupled multiprocessor.

5 A distributed shared memory multiprocessor ASURA: memory and cache architecture
S. Mori, H. Saito, M. Goshima, S. Tomita, M. Yanagihara, T. Tanaka, D. Fraser, K. Joe, H. Nitta
December 1993 **Proceedings of the 1993 ACM/IEEE conference on Supercomputing**

Full text available: pdf(1.17 MB) Additional Information: full citation, references, citings, index terms

6 Synchronization with multiprocessor caches
Joonwon Lee, Umakishore Ramachandran
May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3

Full text available: pdf(1.18 MB) Additional Information: full citation, abstract, references, citings, index terms

Introducing private caches in bus-based shared memory multiprocessors leads to the cache consistency problem since there may be multiple copies of shared data. However, the ability to snoop on the bus coupled with the fast broadcast capability allows the design of special hardware support for synchronization. We present a new lock-based cache scheme which incorporates synchronization into the cache coherency mechanism. With this scheme high-level synchronization primitives as well as low-le ...

7 Delayed consistency and its effects on the miss rate of parallel programs
Michel Dubois, Jin Chin Wang, Luiz A. Barroso, Kangwoo Lee, Yung-Syau Chen
August 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing**

Full text available: pdf(1.01 MB) Additional Information: full citation, references, citings, index terms

8 Cache memory performance in a unix enviroment
Cedell Alexander, William Keshlear, Furrokh Cooper, Faye Briggs
June 1986 **ACM SIGARCH Computer Architecture News**, Volume 14 Issue 3

Full text available: pdf(2.10 MB) Additional Information: full citation, citings, index terms

9 An economical solution to the cache coherence problem
James Archibald, Jean Loup Baer
January 1984 **ACM SIGARCH Computer Architecture News , Proceedings of the 11th annual international symposium on Computer architecture**, Volume 12 Issue 3

Full text available: pdf(728.73 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper we review and qualitatively evaluate schemes to maintain cache coherence in tightly-coupled multiprocessor systems. This leads us to propose a more economical (hardware-wise), expandable and modular variation of the "global directory" approach. Protocols for this solution are described. Performance evaluation studies indicate the limits (number of processors, level of sharing) within which this approach is viable.

10 Cache coherence protocols: evaluation using a multiprocessor simulation model
James Archibald, Jean-Loup Baer
September 1986 **ACM Transactions on Computer Systems (TOCS)**, Volume 4 Issue 4

Full text available: pdf(1.79 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Using simulation, we examine the efficiency of several distributed, hardware-based solutions to the cache coherence problem in shared-bus multiprocessors. For each of the approaches, the associated protocol is outlined. The simulation model is described, and results from that model are presented. The magnitude of the potential performance difference between the various approaches indicates that the choice of coherence solution is very important in the design of an efficient shared-bus multi ...

11 Cache coherence in systems with parallel communication channels & many processors
John C. Willis, Arthur C. Sanderson, Charles R. Hill
November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**

Full text available: pdf(868.59 KB) Additional Information: full citation, abstract, references

This paper describes and analyzes two algorithms for maintaining cache coherence in multiprocessor systems with parallel communication channels and many processors. A distributed link-list relates all cache frames representing the same main memory block. Messages traverse the list to maintain list integrity, exclusive ownership, and consistent values. Memory access semantics are equivalent to a shared memory system without caches. Reference latency, efficiency of memory use, and hardware complex ...

[12] Multi-level shared caching techniques for scalability in VMP-M/C

D. R. Cheriton, H. A. Goosen, P. D. Boyle

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture**, Volume 17 Issue 3

Full text available: pdf(1.27 MB) Additional Information: full citation, abstract, references, citings, index terms

The problem of building a scalable shared memory multiprocessor can be reduced to that of building a scalable memory hierarchy, assuming interprocessor communication is handled by the memory system. In this paper, we describe the VMP-MC design, a distributed parallel multi-computer based on the VMP multiprocessor design, that is intended to provide a set of building blocks for configuring machines from one to several thousand processors. VMP-MC uses a memory hierarchy based on shared caches ...

[13] Options for dynamic address translation in COMAs

Xiaogang Qiu, Michel Dubois

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture**, Volume 26 Issue 3

Full text available: pdf(1.37 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

In modern processors, the dynamic translation of virtual addresses to support virtual memory is done before or in parallel with the first-level cache access. As processor technology improves at a rapid pace and the working sets of new applications grow insatiably the latency and bandwidth demands on the TLB (Translation Lookaside Buffer) are getting more and more difficult to meet. The situation is worse in multiprocessor systems, which run larger applications and are plagued by the TLB consiste ...

[14] A class of compatible cache consistency protocols and their support by the IEEE futurebus

P. Sweazey, A. J. Smith

June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture**, Volume 14 Issue 2

Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, references, citings, index terms

Standardization of a high performance blackplane bus, so that it can accommodate boards developed by different vendors, implies the need for a standardized cache consistency protocol. In this paper we define a class of compatible consistency protocols supported by the current IEEE Futurebus design. We refer to this class as the MOESI class of protocols; the term "MOESI" is derived from the names of the states. This class of protocols has the property that any system component ca ...

[15] Verification of an Industrial CC-NUMA Server

Rajarshi Mukherjee, Yozo Nakayama, Toshiya Mima

January 2002 **Proceedings of the 2002 conference on Asia South Pacific design automation/VLSI Design**

Full text available: pdf(141.86 KB) Publisher Site Additional Information: full citation, abstract

Directed tes program-based verification or formal verification methods are usually quite ineffective on large cache-coherent, non-uniform memory access (CC-NUMA) multi-processors because of the size and complexity of the design and the complexity of the cache-coherence protocol. A controllable biased/constrained random stimuli generator coupled with an error detection mechanism using scoreboards and feedback with coverage analysis tools is a promising alternative methodology. We applied this met ...

[16] Hive: fault containment for shared-memory multiprocessors

J. Chapin, M. Rosenblum, S. Devine, T. Lahiri, D. Teodosiu, A. Gupta

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available: pdf(1.90 MB) Additional Information: full citation, references, citings, index terms

[17] Owner prediction for accelerating cache-to-cache transfer misses in a cc-NUMA architecture

Manuel E. Acacio, José González, José M. García, José Duato

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available: pdf(120.57 KB) Additional Information: full citation, abstract, references, index terms

Cache misses for which data must be obtained from a remote cache (cache-to-cache transfer misses)

account for an important fraction of the total miss rate. Unfortunately, cc-NUMA designs put the access to the directory information into the critical path of 3-hop misses, which significantly penalizes them compared to SMP designs. This work studies the use of owner prediction as a means of providing cc-NUMA multiprocessors with a more efficient support for cache-to-cache transfer misses. Our propo ...

[18] A cache coherence approach for large multiprocessor systems

J. K. Archibald

June 1988 **Proceedings of the 2nd international conference on Supercomputing**

Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper explores the architecture of high-performance large scale multiprocessors using private caches for each processor. The caches reduce the average memory access time, but they also result in the well known cache coherence problem. Multiple copies of each memory location are allowed to exist but they must be kept consistent with each other. In this paper, we present a solution to the cache coherence problem specifically for shared bus multiprocessors that adapts dyn ...

[19] Token coherence: decoupling performance and correctness

Milo M. K. Martin, Mark D. Hill, David A. Wood

May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2

Full text available: pdf(269.08 KB) Additional Information: full citation, abstract, references

Many future shared-memory multiprocessor servers will both target commercial workloads and use highly-integrated "glueless" designs. Implementing low-latency cache coherence in these systems is difficult, because traditional approaches either add indirection for common cache-to-cache misses (directory protocols) or require a totally-ordered interconnect (traditional snooping protocols). Unfortunately, totally-ordered interconnects are difficult to implement in glueless designs. An ideal coherenc ...

[20] Multicast snooping: a new coherence method using a multicast address network

E. Ender Bilir, Ross M. Dickson, Ying Hu, Manoj Plakal, Daniel J. Sorin, Mark D. Hill, David A. Wood

May 1999 **ACM SIGARCH Computer Architecture News , Proceedings of the 26th annual international symposium on Computer architecture**, Volume 27 Issue 2

Full text available: pdf(98.12 KB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

This paper proposes a new coherence method called "multicast snooping" that dynamically adapts between broadcast snooping and a directory protocol. Multicast snooping is unique because processors predict which caches should snoop each coherence transaction by specifying a multicast *"mask."* Transactions are delivered with an ordered multicast network, such as an Isotach network, which eliminates the need for acknowledgment messages. Processors handle transactions as they would with a snoop ...

Terms used
coherency directory and **owner** or **ownership** and **tag** and **owned** and **global**

Found **2,797** of **138,517**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Adjustable block size coherent caches

Czarek Dubnicki, Thomas J. LeBlanc

April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual international symposium on Computer architecture**, Volume 20 Issue 2

Full text available: pdf(1.24 MB) Additional Information: full citation, abstract, references, citings, index terms

Several studies have shown that the performance of coherent caches depends on the relationship between the granularity of sharing and locality exhibited by the program and the cache block size. Large cache blocks exploit processor and spatial locality, but may cause unnecessary cache invalidations due to false sharing. Small cache blocks can reduce the number of cache invalidations, but increase the nuber of bus or network transactions required to load data into the cache. In this paper we ...

2 Cache memory performance in a unix enviroment

Cedell Alexander, William Keshlear, Furrokh Cooper, Faye Briggs

June 1986 **ACM SIGARCH Computer Architecture News**, Volume 14 Issue 3

Full text available: pdf(2.10 MB) Additional Information: full citation, citings, index terms

3 Memory coherence in shared virtual memory systems

Kai Li, Paul Hudak

November 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 4

Full text available: pdf(2.71 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The memory coherence problem in designing and implementing a shared virtual memory on loosely coupled multiprocessors is studied in depth. Two classes of algorithms, centralized and distributed, for solving the problem are presented. A prototype shared virtual memory on an Apollo ring based on these algorithms has been implemented. Both theoretical and practical results show that the memory coherence problem can indeed be solved efficiently on a loosely coupled multiprocessor.

4 Delayed consistency and its effects on the miss rate of parallel programs

Michel Dubois, Jin Chin Wang, Luiz A. Barroso, Kangwoo Lee, Yung-Syau Chen

August 1991 **Proceedings of the 1991 ACM/IEEE conference on Supercomputing**

Full text available: pdf(1.01 MB) Additional Information: full citation, references, citings, index terms

5 A class of compatible cache consistency protocols and their support by the IEEE futurebus

P. Sweazey, A. J. Smith

June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture**, Volume 14 Issue 2

Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, references, citings, index terms

Standardization of a high performance blackplane bus, so that it can accommodate boards developed by different vendors, implies the need for a standardized cache consistency protocol. In this paper we define a class of compatible consistency protocols supported by the current IEEE Futurebus design. We refer to this class as the MOESI class of protocols; the term "MOESI" is derived from the names of the states. This class of protocols has the property that any system component ca ...

6 Multi-level shared caching techniques for scalability in VMP-M/C

D. R. Cheriton, H. A. Goosen, P. D. Boyle

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture**, Volume 17 Issue 3

Full text available: pdf(1.27 MB) Additional Information: full citation, abstract, references, citings, index terms

The problem of building a scalable shared memory multiprocessor can be reduced to that of building a scalable memory hierarchy, assuming interprocessor communication is handled by the memory system. In this paper, we describe the VMP-MC design, a distributed parallel multi-computer based on the VMP multiprocessor design, that is intended to provide a set of building blocks for configuring machines from one to several thousand processors. VMP-MC uses a memory hierarchy based on shared caches ...

7 Cache coherence protocols: evaluation using a multiprocessor simulation model

James Archibald, Jean-Loup Baer

September 1986 **ACM Transactions on Computer Systems (TOCS)**, Volume 4 Issue 4

Full text available: pdf(1.79 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Using simulation, we examine the efficiency of several distributed, hardware-based solutions to the cache coherence problem in shared-bus multiprocessors. For each of the approaches, the associated protocol is outlined. The simulation model is described, and results from that model are presented. The magnitude of the potential performance difference between the various approaches indicates that the choice of coherence solution is very important in the design of an efficient shared-bus multi ...

8 A cache consistency protocol for multiprocessors with multistage networks

P. Stenström

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture**, Volume 17 Issue 3

Full text available: pdf(920.91 KB) Additional Information: full citation, abstract, references, citings, index terms

A hardware based cache consistency protocol for multiprocessors with multistage networks is proposed. Consistency traffic is restricted to the set of caches which have a copy of a shared block. State information is distributed to the caches and the memory modules need not be consulted for consistency actions. The protocol provides two operating modes: distributed write and global read. Distribution of writes calls for efficient multicast methods. Communication cost for multicasti ...

9 A cache coherence approach for large multiprocessor systems

J. K. Archibald

June 1988 **Proceedings of the 2nd international conference on Supercomputing**

Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper explores the architecture of high-performance large scale multiprocessors using private caches for each processor. The caches reduce the average memory access time, but they also result in the well known cache coherence problem. Multiple copies of each memory location are allowed to exist but they must be kept consistent with each other. In this paper, we present a solution to the cache coherence problem specifically for shared bus multiprocessors that adapts dyn ...

10 The sun fireplane system interconnect

Alan Charlesworth

November 2001 **Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(224.87 KB) Additional Information: full citation, abstract, references, citings, index terms

System interconnect is a key determiner of the cost, performance, and reliability of large cache-coherent, shared-memory multiprocessors. Interconnect implementations have to accommodate ever greater numbers of ever faster processors. This paper describes the Sun™ Fireplane two-level cache-coherency protocol, and its use in the medium and large-sized UltraSPARC-III-based Sun Fire™ servers.

11 Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren

November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28 , 34 Issue 5 , 5

Full text available: pdf(413.91 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an

aggressive IO subsystem. The current implementation supports up to 8 such nodes for a total of 32 processors. While s ...

[12] Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: pdf(1.67 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-coherent non-uniform memory access multiprocessor developed at Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiprocessing with 32 to 64 processors. Each node in the design consists of four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. The current implementation supports up to 8 such nodes for a total of 32 processors. While s ...

[13] Implementing a cache consistency protocol

R. H. Katz, S. J. Eggers, D. A. Wood, C. L. Perkins, R. G. Sheldon

June 1985 **ACM SIGARCH Computer Architecture News , Proceedings of the 12th annual international symposium on Computer architecture**, Volume 13 Issue 3

Full text available: pdf(803.11 KB) Additional Information: full citation, citings, index terms

Keywords: ownership-based protocols, shared bus multicomprocessor cache consistency, single chip implementation, snooping caches

[14] Sensor databases: Cache-and-query for wide area sensor databases

Amol Deshpande, Suman Nath, Phillip B. Gibbons, Srinivasan Seshan

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on on Management of data**

Full text available: pdf(230.75 KB) Additional Information: full citation, abstract, references, index terms

Webcams, microphones, pressure gauges and other sensors provide exciting new opportunities for querying and monitoring the physical world. In this paper we focus on querying *wide area sensor databases*, containing (XML) data derived from sensors spread over tens to thousands of miles. We present the first scalable system for executing XPATH queries on such databases. The system maintains the logical view of the data as a single XML document, while physically the data is fragmented across a ...

[15] Piranha: a scalable architecture based on single-chip multiprocessing

Luiz André Barroso, Kourosh Gharachorloo, Robert McNamara, Andreas Nowatzyk, Shaz Qadeer, Barton Sano, Scott Smith, Robert Stets, Ben Verghese

May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th annual international symposium on Computer architecture**, Volume 28 Issue 2

Full text available: pdf(191.10 KB) Additional Information: full citation, abstract, references, citings, index terms

The microprocessor industry is currently struggling with higher development costs and longer design times that arise from exceedingly complex processors that are pushing the limits of instruction-level parallelism. Meanwhile, such designs are especially ill suited for important commercial applications, such as on-line transaction processing (OLTP), which suffer from large memory stall times and exhibit little instruction-level parallelism. Given that commercial applications constitute by fa ...

[16] A distributed shared memory multiprocessor ASURA: memory and cache architecture

S. Mori, H. Saito, M. Goshima, S. Tomita, M. Yanagihara, T. Tanaka, D. Fraser, K. Joe, H. Nitta

December 1993 **Proceedings of the 1993 ACM/IEEE conference on Supercomputing**

Full text available: pdf(1.17 MB) Additional Information: full citation, references, citings, index terms

[17] An empirical evaluation of two memory-efficient directory methods

Brian W. O'Krafka, A. Richard Newton

May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3

Full text available: pdf(1.19 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents an empirical evaluation of two memory-efficient directory methods for maintaining coherent caches in large shared memory multiprocessors. Both directory methods are modifications of a scheme proposed by Censier and Feautrier [5] that does not rely on a specific

interconnection network and can be readily distributed across interleaved main memory. The schemes considered here overcome the large amount of memory required for tags in the original scheme in two different ways ...

18 Pre-silicon verification of the Alpha 21364 microprocessor error handling system

Richard Lee, Benjamin Tsien

June 2001 **Proceedings of the 38th conference on Design automation**

Full text available: pdf(176.87 KB) Additional Information: full citation, abstract, references, index terms

This paper presents the strategy used to verify the error logic in the Alpha 21364 microprocessor. Traditional pre-silicon strategies of focused testing or unit-level random testing yield limited results in finding complex bugs in the error handling logic of a microprocessor. This paper introduces a technique to simulate error conditions and their recovery in a global environment using random test stimulus closely approximating traffic found in a real system. A significant number of bugs ...

19 Parallel architectures: Inferential queueing and speculative push for reducing critical communication latencies

Ravi Rajwar, Alain Kägi, James R. Goodman

June 2003 **Proceedings of the 17th annual international conference on Supercomputing**

Full text available: pdf(568.93 KB) Additional Information: full citation, abstract, references, index terms

Communication latencies within critical sections constitute a major bottleneck in some classes of emerging parallel workloads. In this paper, we argue for the use of Inferentially Queued Locks (IQLs) [31], not just for efficient synchronization but also for reducing communication latencies, and we propose a novel mechanism, Speculative Push (SP), aimed at reducing these communication latencies. With IQLs, the processor infers the existence, and limits, of a critical section from the use of synch ...

Keywords: data forwarding, inferential queueing, synchronization

20 Options for dynamic address translation in COMAs

Xiaogang Qiu, Michel Dubois

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture**, Volume 26 Issue 3

Full text available: pdf(1.37 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

In modern processors, the dynamic translation of virtual addresses to support virtual memory is done before or in parallel with the first-level cache access. As processor technology improves at a rapid pace and the working sets of new applications grow insatiably the latency and bandwidth demands on the TLB (Translation Lookaside Buffer) are getting more and more difficult to meet. The situation is worse in multiprocessor systems, which run larger applications and are plagued by the TLB consiste ...

Nothing Found

Your search for **"coherency directory" or "coherence directory" or "global directory"** did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

- Enter your search terms in lower case with a space between the terms.

  ```
  sales offices
  ```

  You can also enter a full question or concept in plain language.

  ```
  Where are the sales offices?
  ```

- Capitalize proper nouns to search for specific people, places, or products.

  ```
  John Colter, Netscape Navigator
  ```

- Enclose a phrase in double quotes to search for that exact phrase.

  ```
  "museum of natural history" "museum of modern art"
  ```

- Narrow your searches by using a + if a search term must appear on a page.

  ```
  museum +art
  ```

- Exclude pages by using a - if a search term must not appear on a page.

  ```
  museum -Paris
  ```

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

```
museum +"natural history" dinosaur -Chicago
```